

WP4 Status and Plans - Operations Optimizations


Jack C. Wells (for OLCF Team and BigPanDA@Titan)
Director of Science
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

BigPanDA@OLCF Technical Interchange Meeting

University of Texas at Arlington

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. Some of the work presented here is from the TOTAL and Oak Ridge National Laboratory collaboration which is done under the CRADA agreement NFE-14-05227. Some of the experiments were supported by an allocation of advanced computing resources provided by the National Science Foundation. The computations were performed on Nautilus at the National Institute for Computational Sciences.


ORNL is managed by UT-Battelle
for the US Department of Energy



Track #4: Monitoring shared services

- Quantify the compute hours delivered through PanDA that would not otherwise been possible to deliver.
- Monitor the scheduling and delivery of leadership-class jobs on Titan during the operation of the PanDA workflow.
- Monitor the impact of the I/O associated with the PanDA-enabled payloads on the wider Titan I/O performance.
- Quantify the marginal increase in network contention on Titan's interconnection network.
- Monitor the marginal increase in contention for access to the data transfer nodes (DTN).
- Monitor the increased load on the service-node infrastructure due to the PanDA Pilot jobs.
- Quantify the marginal increase in WAN traffic into and out of the OLCF data center.


Titan load (%) for Q4 CY 2017



- Frequent full machine “drains” are consistent with end-of-the-year, full-scale INCITE projection jobs.

Operational Highlight: Making Use of Backfill

- Consumed 173 Million Titan core hours in 2017
 - This is 3.6 percent of total available time on Titan in 2017
 - 230% of average INCITE project
- 8.3% of Titan remained unused



CSC108 (ATLAS backfill) fits comfortably within overall OLCF Discretionary Program without disrupting other DD projects


CSC108 consumption as part of OLCF DD Program

173 M (CSC108)

+379 M (All other DD Projects)

552 M (Total DD in 2017)

✓ 375 M is the total DD program target for 2017



2017 OLCF Operational Year End Metrics (draft totals)

- Over-delivered on our INCITE and ALCC core-hour commitments for the year
 - INCITE at 60.37% of total delivered
 - ALCC at 26.99% of total delivered
 - DD at 12.63% of total delivered
- Delivered a record number of hours to our user programs with 4.82 billion core-hours available and 4.40 billion consumed
 - These are both up from last year's numbers of 4.77 and 4.32 billion, respectively
- We reduced the number of downtimes from 17 in 2016 to 12 in 2017 and only had 4 unscheduled outages
 - Overall Availability increased from 97% in 2016 to 98% in 2017
- Our capability metrics also increased for 2017
 - INCITE capability usage increased from 62% in 2016 to 68% in 2017
 - Overall capability usage went from 59% in 2016 to 60% in 2017

Solving of IO crisis with ATLAS production at OLCF

Danila Oleynik UTA

Profiling Athena-MP IO on Lustre


- Attempts to profile IO on Spider (Lustre) were frustrated by inability to compile Athena with the Darshan tool.
- However, Sarp's analysis of the Lustre logs gave indication that an abundance of file metadata operations could be a weakness and a problem.

	Num. Nodes	Duration (s)	Read (GB)	Written (GB)	GB Read/nodes	GB Written/nodes	<i>open()</i>	<i>close()</i>
Min	1	1,932	0.01	0.03	0.00037	0.02485	1,368	349
Max	300	7,452	241.06	71.71	0.81670	0.23903	1,260,185	294,908
Average	35.66	6,280.82	20.36	6.87	0.38354	0.16794	146,459.37	34,155.74
Std. Dev.	55.33	520.99	43.90	12.33	0.19379	0.03376	231,346.55	53,799.08

Table 1: The Statistical breakdown of the I/O impact of 1,175 jobs ATLAS executed at OLCF for the week of 10/25/16

Unpublished

Moving of working directory to RAM-disk. Results



- Significant reduction of IO. Number of 'open' operations almost matches with number of close operations.
 - Initial spike came from MPI wrapper which used to launch ATHENA Job on computing node. Already reduced with same fix like athena.py
- Current setup of ATLAS production at OLCF
 - ATLAS releases: NFS
 - Job working directories and input data: RAM disk of computing node
 - Output data moved to Lustre at the end of the job

Ack: Danila Oleynik

IO optimization of ATLAS production on OLCF Spider file system

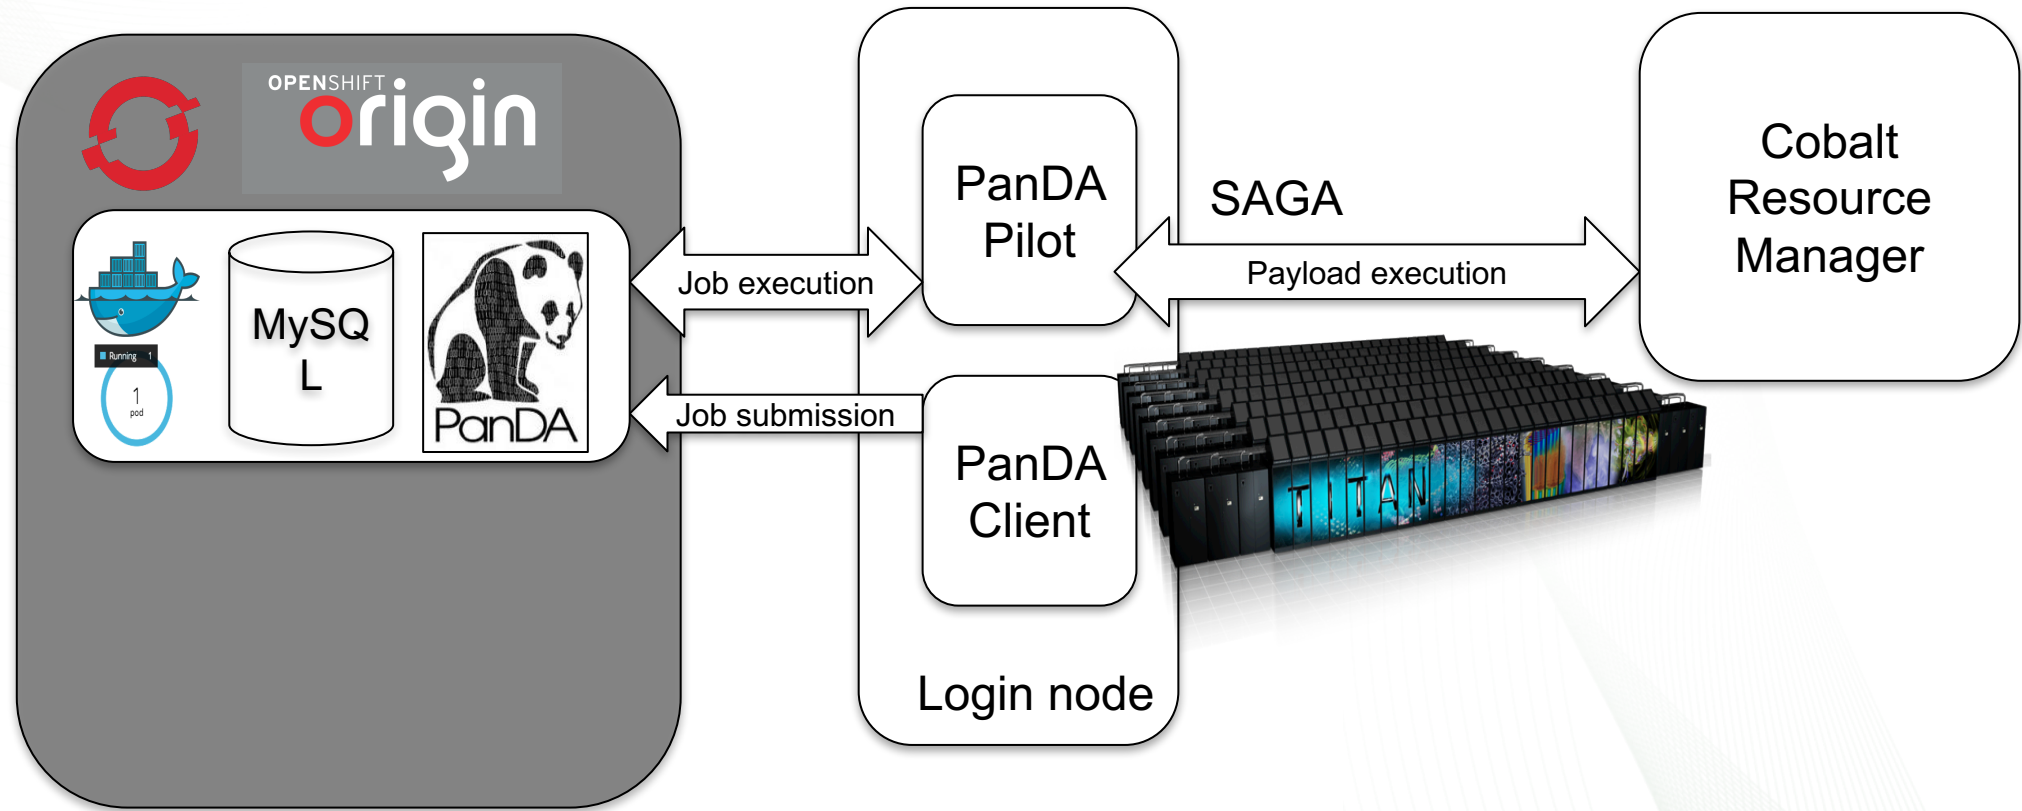
- It took about 2 month of investigation, fixing and proving of solution
- Different specialist were involved:
 - OLCF support team and sysadmins,
 - ATHENA developers,
 - PanDA team
- This is a site-specific optimization, but lessons learned can be used by other site and future development

Ack: Danila Oleynik

Track #4: Monitoring shared services

- Quantify the compute hours delivered through PanDA that would not otherwise been possible to deliver.
- Monitor the scheduling and delivery of leadership-class jobs on Titan during the operation of the PanDA workflow.
- Monitor the impact of the I/O associated with the PanDA-enabled payloads on the wider Titan I/O performance.
- Quantify the marginal increase in network contention on Titan's interconnection network.
- Monitor the marginal increase in contention for access to the data transfer nodes (DTN).
- Monitor the increased load on the service-node infrastructure due to the PanDA Pilot jobs.
- Quantify the marginal increase in WAN traffic into and out of the OLCF data center.

PanDA service under OpenShift Origin at ORNL

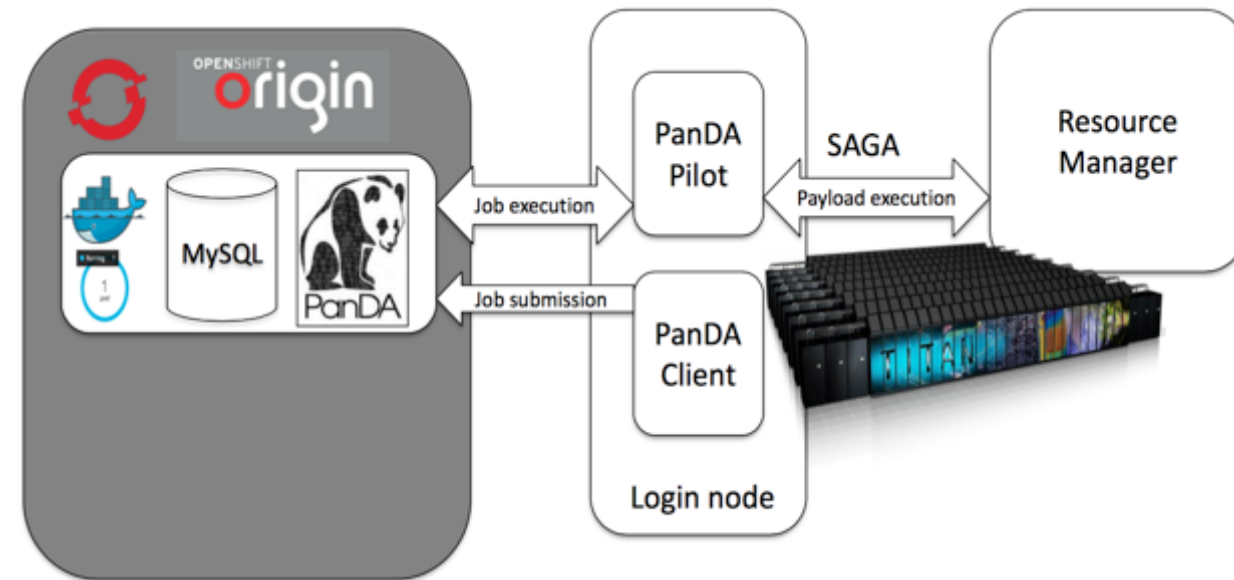



PanDA Server at OLCF: Broad application across domains

- In March 2017 a new PanDA server instance has been set up at ORNL to serve various experiments. This installation is using a container cluster management and orchestration system, Red Hat OpenShift Origin.
 - We are looking forward for experimenting with GPU payloads with PanDA provided by ATLAS b-tagging team and BER projects.
- New experiments to be served by the new PanDA Server instance



- Experiments currently evaluating PanDA



Key Collaborators:

Ruslan.Mashinistov (U. Texas – Arlington)

Jason Kincl (ORNL)

PanDA Server Instance @ OLCF: This meeting

- Day 1, 16:00 - 16:15, BigPanDA Server at OLCF
 - Speaker: Ruslan Mashinistov (University of Texas at Arlington)
- Day 1, 16:15 – 18:15, 8 Experiments/User Talks
- Day 2, 11:00 – 12:00, 2018 Planning Discussion

PanDA Server Instance @ OLCF: Planning for Operations

- Envision two operational modes (at minimum).
 - (1) Test bed/playground for novel or new user projects
 - (2) Production capabilities for science user projects who have adopted PanDA
- Leadership:
 - (1) NCCS/OLCF Advanced Data & Workflow Group (ADW, A. Shankar)
 - (2) Production user project will be in the lead
- Support:
 - HPC Ops envisions supporting this through OpenSHIFT
 - Long-term relationship with PanDA team needs to be envisioned.

Future Plans

- Continue CSC108 ATLAS DD project – scale up backfill
- Compare/contrast CSC108 with ALCC ATLAS project (PI, Childers)
- Support implementation of Harvester edge services
- Scale up diverse science projects using PanDA Instance @ OLCF
 - Detailed summaries on Day 2.
- Transfer PanDA operations to NCCS/OLCF ADW Group & additional Science projects.

Questions